

DECOMPOSITION OF LARGE-SCALE STOCHASTIC OPTIMAL CONTROL PROBLEMS

KENGY BARTY, PIERRE CARPENTIER, AND PIERRE GIRARDEAU

ABSTRACT. In this paper, we present an Uzawa-based heuristic that is adapted to some type of stochastic optimal control problems. More precisely, we consider dynamical systems that can be divided into small-scale independent subsystems, though linked through a static almost sure coupling constraint at each time step. This type of problem is common in production/portfolio management where subsystems are, for instance, power units, and one has to supply a stochastic power demand at each time step. We outline the framework of our approach and present promising numerical results on a simplified power management problem.

1. INTRODUCTION

Stochastic optimal control is concerned about finding strategies to manage dynamical systems in an optimal way, with respect to some cost function. The particularity of such optimization problems is that the optimization variables we deal with are random variables. Indeed, the dynamical systems we consider are partially driven by some exogeneous noises and the objective function may also include such noises. Hence controls are random variables. Classical approaches such as Dynamic Programming and Stochastic Programming, that we briefly recall below, encounter difficulties when the system becomes large. The aim of this paper is to present a new heuristic to solve a class of such problems using price decomposition.

The problems we are studying are common in practice. For example, consider a physical system, say a set of numerous power units, that evolve depending on exogeneous noises (water inflows, failures) and on controls (production levels). At each time step, an observation on the system arises and a control has to be chosen on the basis of the available information, namely the past observations (non-anticipativity constraint). The objective is to minimize the sum of the units' production costs over a given discretized time horizon, while satisfying a global demand constraint at each time step. This decision process hence consists of finding optimal strategies, i.e. functions that map, at each time t , the available information to the optimal decision with respect to the production cost.

As far as we know, most methods that have been proposed to decompose large-scale stochastic optimal control problems are based on Stochastic Programming (see [Pré95, SR03]). This approach consists in representing the non-anticipativity constraints using a so-called scenario tree. Once discretized on such a structure, the problem is not stochastic anymore and various deterministic decomposition

Date: March 6, 2009.

1991 Mathematics Subject Classification. 93E20, 49M27, 49L20.

Key words and phrases. Stochastic optimal control, Decomposition methods, Dynamic programming.

techniques have been used to solve it (see [HS96]). In this context, there are two main issues that are not easy to deal with. The first is concerned with the “distance” between the original problem and its deterministic reformulation [Art91, HRS06], or how to draw a scenario tree in such a way that the solution of the discretized problem is an accurate estimate of the original one. In order to obtain some given accuracy, [Sha06] shows that the growth of the number of leaves in the tree, hence the numerical complexity, has to be exponential with respect to the time horizon. The second issue is concerned with the way one can rebuild strategies from optimal commands obtained in the discretized problem [Pen05].

On the other hand, when dealing with a Markov Decision Process, methods based on Dynamic Programming (DP) (see [Bel57, Ber00]) do provide a way to obtain strategies as feedback functions with respect to so-called state variables. Unfortunately, the well-known curse of dimensionality prevents us from using this approach straightforward on large-scale problems, because the computational burden increases exponentially with the state dimension. Numerous approximations have been proposed to tackle the difficulty. For instance, a popular idea in the field of hydro-power management, introduced by Turgeon in [Tur80], consists of obtaining local strategies as a function of the local stock and the aggregated complementary stock. Another idea, namely Approximate Dynamic Programming (ADP), looks for the value functions (solutions of the DP equation) within a finite-dimensional space (see [BD59] or [BT96, §6.5]). To be practically efficient, such an approach requires some a priori information about the problem, in order to define a well suited functional subspace. Indeed, there is no systematic means to choose the basis functions and several choices have been proposed [TVR96, dFVR03].

When dealing with large-scale optimization problems, the decomposition/coordination approach aims at finding a solution to the original problem by iteratively solving several smaller-dimensional subproblems. In the deterministic case, several types of decomposition have been proposed (e.g. by prices or by quantities) and unified in a general framework using the Auxiliary Problem Principle in [Coh80]. In the open-loop stochastic case, i.e. when controls do not rely on any observation, [CC90] proposed to take advantage of both decomposition techniques and stochastic gradient algorithms. These techniques have been extended in the closed-loop stochastic case by [BRS07], but so far they fail to provide decomposed state dependent strategies in the Markovian case. This is because a subproblem’s optimal strategy depends on the state of the whole system, not only on the local state. In other words, decomposition approaches are meant to decompose the control space, namely the range of the strategy, but the numerical complexity of the problems we consider here also arises because of the dimensionality of the state space, that is to say the domain of the strategy.

We here propose a way to use price decomposition within the closed-loop stochastic case. The coupling constraints, namely the constraints preventing the problem from being naturally decomposed, are dualized using a Lagrange multiplier (price). At each iteration, the price decomposition algorithm solves each subproblem using the current prices, then uses the solutions to update the prices. In the stochastic context, prices are a random process whose dynamics are not available, so the subproblems do not in general fall into the Markovian setting. However, on a specific instance of this problem, [Str06] has exhibited a dynamics for the optimal multipliers, and he has shown that these dynamics were independent of the

decision variables. Hence it was possible to come down to the Markovian framework and to use DP to solve the subproblems in this case. Following this idea, we propose to choose a parameterized dynamics for these multipliers so that solving subproblems using DP becomes possible. The update is then performed using a sampling/regression technique.

This paper is organized as follows. In §2 we describe the general type of problems we are concerned with in this paper. Then, in §3, we recall the Dynamic Programming equation and highlight the difficulties induced when considering large-scale problems. In §4, we present the classical price decomposition approach in Hilbert spaces and the difficulties encountered when dealing with stochastic optimal control problems. Based on these ingredients, we present in §5 a heuristic allowing us to solve subproblems using DP. We finally validate this approach on a simplified power management problem in §6.

2. MATHEMATICAL FRAMEWORK

Throughout this paper the random variables, defined over a probability space $(\Omega, \mathcal{A}, \mathbb{P})$, will be denoted using bold letters (e.g. $\mathbf{W} \in L^2(\Omega, \mathcal{A}, \mathbb{P}, \mathbb{W})$) whereas their realizations will be denoted using normal letters (e.g. $w \in \mathbb{W}$).

In this paper we consider a finite horizon stochastic optimal control problem, where T denotes the time horizon. Three types of random variables are involved in the problem, namely a state, a control, and a noise. The state $\mathbf{X}_t \in L^2(\Omega, \mathcal{A}, \mathbb{P}, \mathbb{R}^n)$ evolves with respect to dynamics depending on the control $\mathbf{U}_t \in L^2(\Omega, \mathcal{A}, \mathbb{P}, \mathbb{R}^m)$ and on some exogeneous noise $\boldsymbol{\xi}_t \in L^2(\Omega, \mathcal{A}, \mathbb{P}, \mathbb{R}^p)$. Unlike deterministic optimal control problems, in the stochastic case the time “direction” is of particular importance. In order to fulfill the causality principle, the control at a given time step t only depends on the observation of noises prior to t . Moreover, we assume that the observation available at time t consists of all past noises. In order to mathematically represent such an information structure, we denote by \mathcal{A}_t the σ -field generated at time t by past noises $(\boldsymbol{\xi}_1, \dots, \boldsymbol{\xi}_t)$, so that the control \mathbf{U}_t at time step t has to be measurable with respect to \mathcal{A}_t . These last constraints will be called the non-anticipativity constraints.

The global system consists of N units, whose dynamics and cost functions are mutually independent. More precisely, the state \mathbf{X}_t (respectively the control \mathbf{U}_t) of the global system writes $(\mathbf{X}_t^1, \dots, \mathbf{X}_t^N)$ with $\mathbf{X}_t^i \in L^2(\Omega, \mathcal{A}, \mathbb{P}, \mathbb{R}^{n_i})$ (respectively $(\mathbf{U}_t^1, \dots, \mathbf{U}_t^N)$ with $\mathbf{U}_t^i \in L^2(\Omega, \mathcal{A}, \mathbb{P}, \mathbb{R}^{m_i})$) and $n = \sum_{i=1}^N n_i$ and $m = \sum_{i=1}^N m_i$, so that the global dynamics $\mathbf{X}_{t+1} = f_t(\mathbf{X}_t, \mathbf{U}_t, \boldsymbol{\xi}_{t+1})$ can be written independently unit by unit: $\mathbf{X}_{t+1}^i = f_t^i(\mathbf{X}_t^i, \mathbf{U}_t^i, \boldsymbol{\xi}_{t+1})$, $i = 1, \dots, N$. In the same way, the global cost $L_t(\mathbf{X}_t, \mathbf{U}_t, \boldsymbol{\xi}_{t+1})$ is equal to the sum of the local unit costs $L_t^i(\mathbf{X}_t^i, \mathbf{U}_t^i, \boldsymbol{\xi}_{t+1})$, $i = 1, \dots, N$. At the end of the time period, each unit i leads a cost K^i that only depends on the final state \mathbf{X}_T^i of the unit.

For now, the global problem can be stated independently unit by unit. The coupling between the units arises from a set of static \mathbb{R}^d -valued constraints, the constraint at time step t reading $\sum_{i=1}^N g_t^i(\mathbf{X}_t^i, \mathbf{U}_t^i) = 0$ (see remark 2 for extensions to more enhanced relations). We suppose that all functions f_t^i , L_t^i and g_t^i are at least Borel measurable.

The initial state \mathbf{X}_0 is assumed to be known. Denoting $(\mathbf{X}_1, \dots, \mathbf{X}_T)$ by \mathbf{X} and $(\mathbf{U}_0, \dots, \mathbf{U}_{T-1})$ by \mathbf{U} , the problem we wish to solve is:

$$\begin{aligned}
 (1a) \quad & \min_{\mathbf{X}, \mathbf{U}} \quad \mathbb{E} \left(\sum_{t=0}^{T-1} \sum_{i=1}^N L_t^i (\mathbf{X}_t^i, \mathbf{U}_t^i, \boldsymbol{\xi}_{t+1}) + \sum_{i=1}^N K^i (\mathbf{X}_T^i) \right) \\
 (1b) \quad & \text{s.t.} \quad \mathbf{X}_{t+1}^i = f_t^i (\mathbf{X}_t^i, \mathbf{U}_t^i, \boldsymbol{\xi}_{t+1}), \quad \forall t = 0, \dots, T-1, \forall i = 1, \dots, N, \\
 (1c) \quad & \mathbf{X}_0^i = x^i, \quad \forall i = 1, \dots, N, \\
 (1d) \quad & \sum_{i=1}^N g_t^i (\mathbf{X}_t^i, \mathbf{U}_t^i) = 0, \quad \forall t = 0, \dots, T-1, \\
 (1e) \quad & \mathbf{U}_t \text{ is } \mathcal{A}_t\text{-measurable}, \quad \forall t = 0, \dots, T-1, \\
 (1f) \quad & \underline{x}_t \leq \mathbf{X}_t \leq \bar{x}_t, \quad \forall t = 1, \dots, T, \\
 (1g) \quad & \underline{u}_t \leq \mathbf{U}_t \leq \bar{u}_t, \quad \forall t = 0, \dots, T-1.
 \end{aligned}$$

There are three types of coupling in Problem (1):

- The first one comes from the state dynamics (1b) that induce a temporal coupling, subsystem by subsystem.
- The second one arises from the static constraints (1d) that link together all the subsystems at each time step t .
- The third type of coupling comes from the non-anticipativity constraints (1e), which link together controls relying on the same noise history. If two realizations of the noise process are identical up to time t , then the same control has to be applied at time t on both realizations.

We ultimately suppose that noises $\boldsymbol{\xi}_t$ are independent (white noise). We are thus in the Markovian case and it is well known that the optimal control, which is a priori a function of all the past noises, only depends on the current state [Ber00].

Remark 1 (White noise assumption). If the noises $\boldsymbol{\xi}_t$ are not independent¹ but still have known dynamics, one can always include the necessary noise history in the state to come back to the Markovian case. Unfortunately, this usually leads to a higher state dimension, and hence a higher numerical complexity in the DP framework, as will be explained in §3.

Remark 2 (Coupling constraints involving noises). It is possible to replace the static coupling constraint $\sum_{i=1}^N g_t^i (\mathbf{X}_t^i, \mathbf{U}_t^i) = 0$ by $\sum_{i=1}^N g_t^i (\mathbf{X}_t^i, \mathbf{U}_t^i) = \mathbf{D}_t$, where \mathbf{D}_t is a random variable representing for instance a global demand. However, expressions are then harder to write: in the Markovian case, i.e. when the \mathbf{D}_t 's are independent one from another, \mathbf{D}_t is observed before choosing the control at time t , so optimal controls must depend on both the state \mathbf{X}_t and the noise \mathbf{D}_t .

3. STOCHASTIC DYNAMIC PROGRAMMING

In order to solve stochastic optimal control problems in the Markovian framework, Bellman proposed in [Bel57] the Dynamic Programming (DP) method. It consists of introducing value functions $V_t : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ that represent the expected

¹and also the noises \mathbf{D}_t introduced in remark 2

optimal cost when starting from state x at time t . In the case of Problem (1), it reads:

$$(2) \quad V_t(x) = \min_{\mathbf{X}, \mathbf{U}} \mathbb{E} \left(\sum_{s=t}^{T-1} \sum_{i=1}^N L_s^i(\mathbf{X}_s^i, \mathbf{U}_s^i, \boldsymbol{\xi}_{s+1}) + \sum_{i=1}^N K^i(\mathbf{X}_T^i) \middle| \mathbf{X}_t = x \right),$$

subject to the same constraints as in (1) and using the convention that if the optimization problem (2) is not feasible, then $V_t(x) = +\infty$. The value functions are usually computed in a recursive manner using the DP equation:

$$(3a) \quad V_T(x) = \sum_{i=1}^N K^i(x^i),$$

and, for $t = 1, \dots, T-1$:

$$(3b) \quad V_t(x) = \min_{u \in [\underline{u}_t, \bar{u}_t]} \mathbb{E} \left(\sum_{i=1}^N L_t^i(x^i, u^i, \boldsymbol{\xi}_{t+1}) + V_{t+1}(f_t(x, u, \boldsymbol{\xi}_{t+1})) \right),$$

$$(3c) \quad \text{s.t. } \sum_{i=1}^N g_t^i(x^i, u^i) = 0.$$

Unlike Stochastic Programming methods, a major advantage of DP is that it provides the control \mathbf{U}_t as a feedback function on the state variable \mathbf{X}_t :

$$\mathbf{U}_t = \Phi_t(\mathbf{X}_t).$$

Except on very simple examples, Equation (3) cannot be solved analytically, and many numerical methods have been proposed. A common practice is to discretize the state space and estimate the expectations using Monte Carlo sampling. Unfortunately, as was mentioned in §1, we are facing the curse of dimensionality: the complexity of DP grows exponentially with respect to the state space dimension.

Moreover, Equation (3) is not decomposable in the sense that it cannot be replaced by the solving of N DP equations depending only on the local state x^i . Indeed, even if V_T is a sum of functions depending on the local state x^i as in (3a), this additive property does not hold for the preceding time steps because of the coupling constraint (3c). Hence, looking for the value function as a sum of functions depending only on the local state would lead to suboptimal strategies. In other words, the local state of a subsystem is not sufficient to take the optimal local decision; some global information about the system is necessary.

Nonetheless, DP remains a seductive approach for small-scale problems since it provides a way to obtain feedback functions. Based on a decomposition scheme presented in §4, we will describe in §5 a heuristic approach in which Problem (1) is decomposed into small-scale subproblems that we solve using DP.

4. PRICE DECOMPOSITION

Let us recall some results about the classical Uzawa algorithm [AHU58], which aims at iteratively getting round the static coupling constraint (1d). When the cost function is additive, this algorithm is also referred to as the price decomposition approach (see [Coh80] for further details). Let us first introduce the Lagrangian of

problem (1):

$$\mathcal{L}(\mathbf{X}, \mathbf{U}, \boldsymbol{\lambda}) = \mathbb{E} \left(\sum_{t=0}^{T-1} \sum_{i=1}^N \left(L_t^i(\mathbf{X}_t^i, \mathbf{U}_t^i, \boldsymbol{\xi}_{t+1}) + \boldsymbol{\lambda}_t^\top g_t^i(\mathbf{X}_t^i, \mathbf{U}_t^i) \right) + \sum_{i=1}^N K^i(\mathbf{X}_T^i) \right),$$

with $\boldsymbol{\lambda}_t \in L^2(\Omega, \mathcal{A}, \mathbb{P}, \mathbb{R}^d)$ the Lagrange multiplier associated to the coupling constraint (1d) and $\boldsymbol{\lambda} = (\boldsymbol{\lambda}_0, \dots, \boldsymbol{\lambda}_{T-1})$. When the Lagrangian has a saddle point, we know from classical duality theory in optimization [ET92] that Problem (1) is equivalent to:

$$(4a) \quad \max_{\boldsymbol{\lambda}} \min_{\mathbf{X}, \mathbf{U}} \mathcal{L}(\mathbf{X}, \mathbf{U}, \boldsymbol{\lambda})$$

$$(4b) \quad \text{s.t. } \mathbf{X}_{t+1}^i = f_t^i(\mathbf{X}_t^i, \mathbf{U}_t^i, \boldsymbol{\xi}_{t+1}), \quad \forall t = 0, \dots, T-1, \forall i = 1, \dots, N,$$

$$(4c) \quad \mathbf{U}_t \text{ is } \mathcal{A}_t\text{-measurable}, \quad \forall t = 0, \dots, T-1,$$

$$(4d) \quad \underline{x}_t \leq \mathbf{X}_t \leq \bar{x}_t, \quad \forall t = 1, \dots, T,$$

$$(4e) \quad \underline{u}_t \leq \mathbf{U}_t \leq \bar{u}_t, \quad \forall t = 0, \dots, T-1.$$

Recall that the Lagrange multiplier $\boldsymbol{\lambda}_t$ can be interpreted as the marginal price one should pay for satisfying the coupling constraint (1d). Because of the \mathcal{A}_t -measurability of the variables involved in this constraint and of the properties of conditional expectation, it is easy to see that we can always choose $\boldsymbol{\lambda}_t$ to be \mathcal{A}_t -measurable.

Let us introduce the dual function $\psi(\boldsymbol{\lambda}) := \min_{\mathbf{X}, \mathbf{U}} \mathcal{L}(\mathbf{X}, \mathbf{U}, \boldsymbol{\lambda})$ subject to constraints (4b), (4c), (4d) and (4e). The key point of the price decomposition algorithm is that computing $\psi(\boldsymbol{\lambda})$ is much easier than solving the original Problem (1). Indeed, one can write:

$$\begin{aligned} \psi(\boldsymbol{\lambda}) &= \min_{\mathbf{X}, \mathbf{U}} \mathbb{E} \left(\sum_{t=0}^{T-1} \sum_{i=1}^N \left(L_t^i(\mathbf{X}_t^i, \mathbf{U}_t^i, \boldsymbol{\xi}_{t+1}) + \boldsymbol{\lambda}_t^\top g_t^i(\mathbf{X}_t^i, \mathbf{U}_t^i) \right) + \sum_{i=1}^N K^i(\mathbf{X}_T^i) \right), \\ &= \sum_{i=1}^N \min_{\mathbf{X}^i, \mathbf{U}^i} \mathbb{E} \left(\sum_{t=0}^{T-1} \left(L_t^i(\mathbf{X}_t^i, \mathbf{U}_t^i, \boldsymbol{\xi}_{t+1}) + \boldsymbol{\lambda}_t^\top g_t^i(\mathbf{X}_t^i, \mathbf{U}_t^i) \right) + K^i(\mathbf{X}_T^i) \right), \end{aligned}$$

so that we replace the solving of an optimization problem with variables (\mathbf{X}, \mathbf{U}) by the solving of N subproblems with variables $(\mathbf{X}^i, \mathbf{U}^i)$.

Given $\boldsymbol{\lambda}^k$, an iteration of the price decomposition algorithm first solves the N subproblems:

$$\begin{aligned} (5a) \quad &\min_{\mathbf{X}^i, \mathbf{U}^i} \mathbb{E} \left(\sum_{t=0}^{T-1} \left(L_t^i(\mathbf{X}_t^i, \mathbf{U}_t^i, \boldsymbol{\xi}_{t+1}) + \boldsymbol{\lambda}_t^{k\top} g_t^i(\mathbf{X}_t^i, \mathbf{U}_t^i) \right) + K^i(\mathbf{X}_T^i) \right) \\ (5b) \quad &\text{s.t. } \mathbf{X}_{t+1}^i = f_t^i(\mathbf{X}_t^i, \mathbf{U}_t^i, \boldsymbol{\xi}_{t+1}), \quad \forall t = 0, \dots, T-1, \\ (5c) \quad &\mathbf{U}_t^i \text{ is } \mathcal{A}_t\text{-measurable}, \quad \forall t = 0, \dots, T-1, \\ (5d) \quad &\underline{x}_t^i \leq \mathbf{X}_t^i \leq \bar{x}_t^i, \quad \forall t = 1, \dots, T, \\ (5e) \quad &\underline{u}_t^i \leq \mathbf{U}_t^i \leq \bar{u}_t^i, \quad \forall t = 0, \dots, T-1. \end{aligned}$$

The Lagrange multiplier $\boldsymbol{\lambda}^k$ is then updated using a gradient-like algorithm. Under standard assumptions², the gradient of ψ is:

$$\nabla_{\boldsymbol{\lambda}_t} \psi \left(\boldsymbol{\lambda}^k \right) = \sum_{i=1}^N g_t^i \left(\mathbf{X}_t^{i,k+1}, \mathbf{U}_t^{i,k+1} \right),$$

where $\mathbf{X}_t^{i,k+1}$ and $\mathbf{U}_t^{i,k+1}$ are the solutions of Problem (5).

At first sight, Problem (5) looks like a stochastic optimal control problem with control \mathbf{U}_t^i and state \mathbf{X}_t^i , the solution of which would be a local feedback on \mathbf{X}_t^i . This contradicts the fact that the solution of Problem (1) is a feedback function on the whole state $(\mathbf{X}_t^1, \dots, \mathbf{X}_t^N)$. In order to understand where this contradiction comes from, one has to highlight the role of $\boldsymbol{\lambda}$ in Problem (5).

5. DUAL APPROXIMATE DYNAMIC PROGRAMMING

Let us take a closer look at Problem (5). First suppose that $\boldsymbol{\lambda}$ is a white noise process. Then Problem (5) lies in the Markovian framework with state \mathbf{X}_t^i and noise $(\boldsymbol{\xi}_t, \boldsymbol{\lambda}_t)$. The optimal control \mathbf{U}_t^i depends only on the local state \mathbf{X}_t^i and one can apply stochastic dynamic programming to solve this small-scale optimal control problem. Unfortunately, we do not know anything about the time correlations of the price process $\boldsymbol{\lambda}$...

Let us now consider the general case. Defining $(\mathbf{X}_t^i, \boldsymbol{\lambda}_1, \dots, \boldsymbol{\lambda}_t)$ as the state at time t , Problem (5) falls in the Markovian setting. In particular, the optimal control \mathbf{U}_t^i is $(\mathbf{X}_t^i, \boldsymbol{\lambda}_1, \dots, \boldsymbol{\lambda}_t)$ -measurable. However DP in this context proves numerically intractable because the state dimension increases with respect to time.

Consider now an intermediate case, and suppose that the dual variable $\boldsymbol{\lambda}$ has a short memory dynamics, for instance that $\boldsymbol{\lambda}_{t+1}$ only depends on $\boldsymbol{\lambda}_t$ and $\boldsymbol{\xi}_{t+1}$:

$$(6) \quad \boldsymbol{\lambda}_{t+1} = h_t(\boldsymbol{\lambda}_t, \boldsymbol{\xi}_{t+1}).$$

Using $(\mathbf{X}_t^i, \boldsymbol{\lambda}_t)$ as the state variable at time t , Problem (5) falls in the Markovian setting. The state dimension does not increase with respect to time anymore and is hopefully small so that Problem (5) can be solved using DP.

In a very specific instance of Problem (1), namely:

$$(7) \quad \begin{aligned} \min_{\mathbf{X}, \mathbf{U}} \quad & \mathbb{E} \left(\sum_{t=0}^{T-1} \sum_{i=1}^N \frac{c_i}{2} (\mathbf{U}_t^i)^2 + \sum_{i=1}^N \frac{\gamma_i}{2} (\mathbf{X}_T^i - \mathbf{X}_0^i)^2 \right) \\ \text{s.t.} \quad & \mathbf{X}_{t+1}^i = \mathbf{X}_t^i - \mathbf{U}_t^i + \mathbf{A}_{t+1}^i, \quad \forall t = 0, \dots, T-1, \\ & \sum_{i=1}^n \mathbf{U}_t^i = \mathbf{D}_t, \quad \forall t = 0, \dots, T-1, \\ & \mathbf{U}_t \text{ is } \mathcal{A}_t\text{-measurable}, \quad \forall t = 0, \dots, T-1, \end{aligned}$$

with $\boldsymbol{\xi}_t = (\mathbf{D}_t, \mathbf{A}_t^1, \dots, \mathbf{A}_t^N)$, [Str06] has brought to light such an intermediate case. Here the dimension of the state \mathbf{X}_t^i (respectively of the control \mathbf{U}_t^i) in the subsystem i is $n_i = 1$ (respectively $m_i = 1$), for $i = 1, \dots, N$. The result is the following.

²See [Dan67] for results on the differentiability of the dual function ψ .

Proposition 1. *If $(\mathbf{D}_t, \mathbf{A}_t^1, \dots, \mathbf{A}_t^N)_{t=0, \dots, T}$ is a white noise process and if there exists $\alpha \in \mathbb{R}^+$ such that $\gamma_i = \alpha c_i, \forall i = 1, \dots, N$, then the optimal Lagrange multipliers satisfy:*

$$\begin{aligned}\boldsymbol{\lambda}_{t+1} &= \boldsymbol{\lambda}_t + \frac{1}{\sum_{i=1}^N \frac{1}{c_i}} \left(\mathbf{D}_{t+1} (1 + \alpha) - \mathbf{D}_t - \alpha \mathbb{E}(\mathbf{D}_{t+1}) \right. \\ &\quad \left. - \alpha (\mathbf{A}_{t+1} - \mathbb{E}(\mathbf{A}_{t+1})) \right), \\ \boldsymbol{\lambda}_0 &= \frac{1}{\sum_{i=1}^N \frac{1}{c_i}} \left(\mathbf{D}_0 (1 - \alpha) - \alpha \sum_{s=1}^T \mathbb{E}(\mathbf{A}_s) - \alpha \sum_{s=1}^{T-1} \mathbb{E}(\mathbf{D}_s) \right).\end{aligned}$$

Using such a dynamics for the multipliers, it is straightforward to show that Problem (7) splits into N independent optimization subproblems. Taking the state variable as $(\mathbf{X}_t^i, \boldsymbol{\lambda}_t, \mathbf{D}_t)$, the i -th subproblem can be solved using DP in dimension 3. In summary, we have replaced one N -dimensional problem by N 3-dimensional problems.

Note that the proportionality assumption on the cost coefficients in proposition 1 is rather unnatural. Nevertheless, it shows that, in some cases, there exist dynamics for the Lagrange multipliers that is independent of the decision variables.

To deal with more general cases, we propose to approximate the dual process $\boldsymbol{\lambda}$ by some parameterized short-memory process. That is, we try to identify the multipliers that are the closest to the optimal ones within a constrained subspace of stochastic processes. This approach is similar to that employed in the Approximate Dynamic Programming (ADP) method. Since it concerns dual variables rather than DP value functions, we refer to this approach as Dual Approximate Dynamic Programming.

The performance of such an approach highly depends on the choice of the subspace of stochastic processes in which we force the multipliers to lie. However, a major difference with ADP techniques is that approximating the dual variables may lead to violations of the coupling constraints. The larger the chosen subspace, the less the coupling constraints will be violated. Moreover, prior information on the problem may be useful to devise a suitable dynamics.

Let us now present the implementation of DADP. We constrain dual variables to satisfy:

$$(8) \quad \boldsymbol{\lambda}_{t+1} = h_{\alpha_t}(\boldsymbol{\lambda}_t, \boldsymbol{\xi}_{t+1}),$$

where h_{α_t} is an a priori chosen function parameterized by $\alpha_t \in \mathbb{R}^q$. We denote by \mathcal{S} the set of all random processes that verify Equation (8) for some real vector $\alpha = (\alpha_1, \dots, \alpha_{T-1})$. Given a vector α^k of coefficients, the first step of DADP is to solve the N subproblems (5) using DP with state $(\mathbf{X}_t^i, \boldsymbol{\lambda}_t)$. In order to update the Lagrange multipliers, we draw s trajectory samples of the noise $\boldsymbol{\xi}$ and integrate the dynamics (5b) and (8) using the optimal feedback laws, thus obtaining s trajectory samples of \mathbf{X}^k , \mathbf{U}^k and $\boldsymbol{\lambda}^k$. We then perform a gradient step on $\boldsymbol{\lambda}$ sample by sample:

$$\boldsymbol{\lambda}_t^{k+\frac{1}{2}, \sigma} = \boldsymbol{\lambda}_t^{k, \sigma} + \rho_t \sum_{i=1}^N g_t^i \left(\mathbf{X}_t^{i, k, \sigma}, \mathbf{U}_t^{i, k, \sigma} \right), \quad \forall \sigma = 1, \dots, s,$$

with ρ_t being well-chosen real values. Finally, we apply a regression operator \mathcal{R}^s on the samples $\boldsymbol{\lambda}^{k+\frac{1}{2}}$ in order to obtain a stochastic process $\boldsymbol{\lambda}^{k+1}$ lying in \mathcal{S} :

$$\begin{aligned}\mathcal{R}^s(\boldsymbol{\lambda}^{k+\frac{1}{2}}) &= \arg \min_{\alpha_0, \dots, \alpha_{T-1}} \sum_{t=0}^{T-1} \sum_{\sigma=1}^s \left\| \boldsymbol{\mu}_t^\sigma - \boldsymbol{\lambda}_t^{k+\frac{1}{2}, \sigma} \right\|_{\mathbb{R}^d}^2 \\ \text{s.t. } \boldsymbol{\mu}_{t+1}^\sigma &= h_{\alpha_t}(\boldsymbol{\mu}_t^\sigma, \boldsymbol{\xi}_{t+1}^\sigma).\end{aligned}$$

This heuristic is outlined in algorithm 1.

Algorithm 1 Dual Approximate Dynamic Programming

Require: $\varepsilon > 0$, $\gamma > 0$, a shape h_α for the prices dynamics, α^0 .

repeat

$k \leftarrow k + 1$

for $i = 1$ to N **do**

Solve i -th subproblem by DP using parameters α^k for the price dynamics, and obtain $\mathbf{X}^{i,k}$ and $\mathbf{U}^{i,k}$. Both implicitly depend on α^k .

end for

Update paramaters α^k :

$$\alpha^{k+1} = \mathcal{R}^s \left(\left(\boldsymbol{\lambda}_t^k + \rho_t \sum_{i=1}^N g_t^i(\mathbf{X}_t^{i,k}, \mathbf{U}_t^{i,k}) \right)_{t=0, \dots, T-1} \right),$$

where:

$$\boldsymbol{\lambda}_{t+1}^k = h_{\alpha_t^k}(\boldsymbol{\lambda}_t^k, \boldsymbol{\xi}_{t+1}), \quad \forall t = 0, \dots, T-1.$$

until $\|\boldsymbol{\lambda}^{k+1} - \boldsymbol{\lambda}^k\| < \varepsilon$

Remark 3 (Convexity of \mathcal{S}). The regression operator \mathcal{R}^s is meant to be a sample-based approximation of the projection operator \mathcal{R} on \mathcal{S} . Since the latter set can be non-convex, $\mathcal{R}(\boldsymbol{\lambda}^{k+\frac{1}{2}})$ is not necessarily unique: this may lead to numerical instabilities.

Remark 4 (Enhancement of \mathcal{S}). It may be desirable to consider a larger set \mathcal{S} in order to estimate more accurately the price process. For instance, one can extend relation (8) in order to include more memory in the process:

$$\boldsymbol{\lambda}_{t+1} = h_{\alpha_t}((\boldsymbol{\lambda}_\tau)_{\tau \leq t}, (\boldsymbol{\xi}_\tau)_{\tau \leq t+1}),$$

However, this will in general increase the numerical complexity of DP in the solving of the subproblems.

Remark 5 (Another formulation). Alternatively, we could have considered a gradient algorithm that iterates directly on the parameters α of the dynamics (8). In this case, since we have no restrictions on α , the feasible set would have been convex. Unfortunately, because the dynamics (8) may be nonlinear with respect to α , the dual function ψ introduced in §4 might be non-concave with respect to α .

6. NUMERICAL EXPERIMENTS

We tested this approach on a simple power management problem. On this small-scale example, we will be able to compare DADP results to those obtained by DP. Consider a power producer who owns two types of power plants:

- Two hydraulic plants that are characterized at each time step t by their water stock \mathbf{X}_t^i and power production \mathbf{U}_t^i , and receive water inflows ξ_{t+1}^i , $i = 1, 2$. These two units are subject to dynamic constraints but are cost-free;
- One thermal unit with a production cost that is quadratic with respect to its production \mathbf{U}_t^3 . There are no dynamics associated with this unit.

Using these plants, the power producer must supply a power demand \mathbf{D}_t at each time step t , over a discrete time horizon of $T = 25$ time steps. All noises, i.e. the demand \mathbf{D}_t and the inflows ξ_t^1 and ξ_t^2 are chosen to be white noise processes.




FIGURE 1. Mean of \mathbf{D}_t , ξ_t^1 and ξ_t^2 over time (in red is one sample trajectory).

In the model we moreover impose small quadratic costs on the hydraulic power productions in order to ensure that, at least in the deterministic framework and without our approximation, the algorithm would build primal iterates that converge

to the optimal solution of the original problem (see [ET92]). The problem reads³:

$$\begin{aligned}
 (9a) \quad & \min_{\mathbf{X}, \mathbf{U}} \mathbb{E} \left(\sum_{t=0}^{T-1} \left(\epsilon (\mathbf{U}_t^1)^2 + \epsilon (\mathbf{U}_t^2)^2 + L_t(\mathbf{U}_t^3) \right) + K^1(\mathbf{X}_T^1) + K^2(\mathbf{X}_T^2) \right) \\
 (9b) \quad & \text{s.t. } \mathbf{X}_{t+1}^i = \mathbf{X}_t^i - \mathbf{U}_t^i + \boldsymbol{\xi}_{t+1}^i, \quad \forall i = 1, 2, \quad \forall t = 0, \dots, T-1, \\
 (9c) \quad & \mathbf{U}_t^1 + \mathbf{U}_t^2 + \mathbf{U}_t^3 = \mathbf{D}_t, \quad \forall t = 0, \dots, T-1, \\
 (9d) \quad & \underline{x}^i \leq \mathbf{X}_t^i \leq \bar{x}^i, \quad \forall i = 1, 2, \quad \forall t = 1, \dots, T, \\
 (9e) \quad & 0 \leq \mathbf{U}_t^i \leq \bar{u}^i, \quad \forall i = 1, 2, \quad \forall t = 0, \dots, T-1, \\
 (9f) \quad & 0 \leq \mathbf{U}_t^3, \quad \forall t = 0, \dots, T-1, \\
 (9g) \quad & \mathbf{U}_t^i \text{ is } \sigma\{\mathbf{D}_0, \boldsymbol{\xi}_0^1, \boldsymbol{\xi}_0^2, \dots, \mathbf{D}_t, \boldsymbol{\xi}_t^1, \boldsymbol{\xi}_t^2\}\text{-measurable}, \quad \forall i = 1, 2, 3.
 \end{aligned}$$

In this problem, the state \mathbf{X}_t is two-dimensional, hence DP remains numerically tractable and we can use the DP solution as a reference. In order to use DADP, we choose an auto-regressive process for the Lagrange multipliers:

$$\begin{aligned}
 (10a) \quad & \boldsymbol{\lambda}_{t+1} = \alpha_t \boldsymbol{\lambda}_t + \beta_t \mathbf{D}_{t+1} + \gamma_t, \\
 (10b) \quad & \boldsymbol{\lambda}_0 = \beta_0 \mathbf{D}_0 + \gamma_0.
 \end{aligned}$$

We then perform the algorithm and depict its convergence in Figure 2. We first draw the values of the dual function ψ introduced in §4 along with iterations (lower curve) and observe that it converges to the optimal value of the original problem computed by DP. Note that each value of ψ is computed by Monte Carlo simulation over 10^3 scenarios. We also draw the cost of the problem with all constraints satisfied (primal cost) at each iteration (upper curve). As explained in §5, DADP does not ensure that the coupling constraint (9c) is satisfied. To circumvent this difficulty, the thermal unit strategy is chosen in the simulation so as to ensure feasibility of the coupling constraint, i.e.:

$$(11) \quad \mathbf{U}_t^3 = \mathbf{D}_t - (\mathbf{U}_t^1 + \mathbf{U}_t^2).$$

That is, DADP returns three strategies, for each of the hydraulic units and for the thermal unit. However, we use relation (11) for the thermal strategy during simulations in order to ensure demand satisfaction.

Figure 2 shows that the algorithm behaves well, in the sense that the value of the objective function converges quite quickly to a neighbourhood of the optimal value. However, even after 100 iterations the curve is still a bit noisy. This is because the price dynamics employed generates a non-convex-set of stochastic processes. Consequently, the least squares problem solved at each iteration is non-linear, and hence small variations in the actual gradient can result in large changes in the calculated gradient.

The key to convergence in DADP is to obtain a dynamics for the Lagrange multipliers that accurately matches the optimal one. We have represented in Figure 3 the dynamics of the multipliers computed by DADP after 10, 20, 50 and 90

³In this example, we consider two hydraulic plants with characteristics:

$$\begin{aligned}
 \underline{x}^1 &= 0, & \bar{x}^1 &= 50, & \bar{u}^1 &= 6, & K^1(x) &= -7x, \\
 \underline{x}^2 &= 0, & \bar{x}^2 &= 40, & \bar{u}^2 &= 6, & K^2(x) &= -12x, & \epsilon &= 0.1
 \end{aligned}$$

where \underline{y} (resp. \bar{y}) denotes a lower (resp. upper) bound for variable y . Moreover, producing u with the thermal plant costs $L_t(u) = u + u^2$.




FIGURE 2. Value of the dual function (blue), of the primal function (red) and optimum (black) along with iterations.

iterations, and those derived from DP. We observe that the approximate price dynamics issued from DADP satisfactorily converges to the optimal one. This




FIGURE 3. Comparison, using 100 samples, of the approximate prices trajectories (red) with the optimal ones (black), after 10, 20, 50, and 90 iterations of the algorithm (from left to right and top to bottom).

indicates that:

- (1) the dual process converged, although the set of stochastic processes defined by (10) is non-convex,
- (2) there is no need to enhance the chosen dynamics (10) for the multipliers in this particular problem.

Note that the optimal prices derived from DP are obtained by numerically differentiating the Bellman functions, hence the numerical instabilities we observe in the lower parts of the DP prices curves.

Remark 6 (Numerical complexity). We chose to validate the method on a two-dimensional power management problem, so we could compute a reference solution using DP. Note that there would be no additional difficulty in implementing the same algorithm with a larger number of hydraulic units, i.e. with a larger-dimensional state. The complexity of DADP grows linearly with the number of subsystems. The most time consuming part of the algorithm is solving the subproblems. However, this calculations can be easily parallelized on a computer, so that the time needed at each DADP iteration remains constant with respect to the number of subsystems.

7. CONCLUSION

In this paper we present an approach, called Dual Approximate Dynamic Programming (DADP), to solve large-scale stochastic optimal control problems in a price decomposition framework, without discretizing randomness. This method employs classical duality theory results to solve decomposable large-scale systems, as it usually is in the deterministic framework. In order to be able to solve subproblems using DP, we suppose that the Lagrange multipliers obey some parameterized dynamics. The DADP algorithm then iterates on the parameters of these dynamics. What is original in this approach is the use of a dual variable in the optimal local feedback functions as an auxiliary variable that sums up the remaining part of the system.

On an example, we show that this approach is very attractive from a numerical point of view. Using rather simple dynamics for the multipliers, we obtained surprisingly good results with a small number of iterations. The main advantage of the method is that the complexity of the algorithm grows linearly with respect to the number of subsystems so that the curse of dimensionality is circumvented for the considered class of problems.

There are still several important theoretical questions. Since we constrain the dual variables to lie in some a priori chosen subset, we cannot state that the coupling constraints will be satisfied. Hence it would be useful to be able to evaluate the distance between the solution given by the heuristic and the feasible set; this would also give clues on how to choose well-suited dynamics for the dual variables on particular problems. Furthermore, the stochastic process subset on which we constrain the dual variables is possibly non-convex. In this context, it might be valuable to use more enhanced numerical methods for the update of the Lagrange multipliers. Further studies will be concerned with these issues.

REFERENCES

- [AHU58] K.J. Arrow, L. Hurwicz, and H. Uzawa, *Studies in linear and nonlinear programming*, Stanford University Press, 1958.

- [Art91] Z. Artstein, *Sensitivity to σ -fields of information in stochastic allocation*, Stochastics and Stochastics Reports **36** (1991), 41–63.
- [BD59] R. Bellman and S.E. Dreyfus, *Functional approximations and dynamic programming*, Math tables and other aides to computation **13** (1959), 247–251.
- [Bel57] R. Bellman, *Dynamic programming*, Princeton University Press, New Jersey, 1957.
- [Ber00] D. P. Bertsekas, *Dynamic programming and optimal control*, 2 ed., vol. 1 & 2, Athena Scientific, 2000.
- [BRS07] K. Barty, J.-S. Roy, and C. Strugarek, *A stochastic gradient type algorithm for closed loop problems*, Mathematical Programming (2007).
- [BT96] D.P. Bertsekas and J.N. Tsitsiklis, *Neuro-Dynamic Programming*, Athena Scientific, 1996.
- [CC90] G. Cohen and J.-C. Culioli, *Decomposition Coordination Algorithms for Stochastic Optimization*, SIAM J. Control Optimization **28** (1990), no. 6, 1372–1403.
- [Coh80] G. Cohen, *Auxiliary Problem Principle and decomposition of optimization problems*, J. Optimization Theory Appl. (1980), no. 32, 277–305.
- [Dan67] J.M. Danskin, *The theory of max-min*, Springer, Berlin, 1967.
- [dFVR03] D.P. de Farias and B. Van Roy, *The Linear Programming Approach to Approximate Dynamic Programming*, Operations Research **51** (2003), no. 6, 850–856.
- [ET92] I. Ekeland and R. Temam, *Convex analysis and variational problems*, Studies in Advanced Mathematics, CRC Press, 1992.
- [HRS06] H. Heitsch, W. Römisch, and C. Strugarek, *Stability of multistage stochastic programs*, SIAM Journal on Optimization **17** (2006), 511–525.
- [HS96] J.L. Higle and S. Sen, *Stochastic decomposition*, Kluwer, Dordrecht, 1996.
- [Pen05] T. Pennanen, *Epi-convergent discretizations of multistage stochastic programs*, Mathematics of Operations Research **30** (2005), 245–256.
- [Pré95] A. Prékopa, *Stochastic programming*, Kluwer, Dordrecht, 1995.
- [Sha06] A. Shapiro, *On complexity of multistage stochastic programs*, Operations Research Letters **34** (2006), 1–8.
- [SR03] A. Shapiro and A. Ruszcynski (eds.), *Stochastic Programming*, Elsevier, Amsterdam, 2003.
- [Str06] C. Strugarek, *Approches variationnelles et autres contributions en optimisation stochastique*, Thèse de doctorat, École Nationale des Ponts et Chaussées, 5 2006.
- [Tur80] A. Turgeon, *Optimal operation of multi-reservoir power systems with stochastic inflows*, Water Resources Research **16** (1980), no. 2, 275–283.
- [TVR96] J.N. Tsitsiklis and B. Van Roy, *Feature-based methods for large scale dynamic programming*, Machine Learning **22** (1996), 59–94.

KENGY BARTY, EDF R&D, 1, AVENUE DU GÉNÉRAL DE GAULLE, 92141 CLAMART CEDEX,
FRANCE

PIERRE CARPENTIER, ÉCOLE NATIONALE SUPÉRIEURE DE TECHNIQUES AVANCÉES (ENSTA),
32, BOULEVARD VICTOR, 75015 PARIS, FRANCE

PIERRE GIRARDEAU, UNIVERSITÉ PARIS-EST, CERMICS, ÉCOLE DES PONTS, CHAMPS SUR
MARNE, 77455 MARNE LA VALLÉE CEDEX 2, FRANCE, ALSO WITH EDF R&D AND ENSTA